

File: TT_paper_thread

From: Tom LaPorta [tlp@dnrc.bell-labs.com]
Sent: Tuesday, January 25, 2000 1:17 PM
To: thuel@lucent.com; ramjee@dnrc.bell-labs.com; lsalgarelli@lucent.com;
kvaradhan@lucent.com; tlp@dnrc.bell-labs.com
Subject: Comments on DHCP paper

Sandy,

Due to the snow, I am trapped in NJ for an extra day and figured I might as well work.

I read the paper - it is very good. Here are my comments (I can give more when I return).

Tom

General

1. You use "should send" and like phrase when you should just state things like "the mobile

sends" or the "the process is halted".

2. Should we call the mobile host a mobile client instead (consistently throughout the paper)?

3. The main benefit of our approach is the we get all the benefits of DHCP (all options, re-use

of servers, etc.). These must be spelled out very clearly or people will say "why assing a 10*

address, just assign the real thing).

4. We need a paragraph on how this would work with external FAs.

5. Title: I think we should name our protocol. My proposal (off the top of my head): Transient


Tunneling Protocol for Dynamic Addressing for Mobile Clients. The TTP could probably find other

uses as well.

Intro

1. Need stong statement of benefits/philosphy of our approach: minor changes to all protocols

and servers, leverage existing s options



Section 4

1. Should figure 3 be broken into an "a" and "b" - one with a rely and one w/o?

Section 5

1. Could use a listing of our benefits
2. Move 3rd paragraph into 5.1.

Section 5.1

1. We don't NEED to tear down the 10* tunnel - soft-state could expire. We should mention this as an option.
2. We should mention procedures with an FA here in 1 of 2 ways - eith put a para here, or point
to a para in the discussion section.
3. Does the B-bits cause any extra broadcast traffic to reach the mobile client/serving subnet?

This may be very undesireable and may be a big drawback of our system.

Section 5.2

1. Is the first para just standard? If yes, say so.

Section 6.

1. Performance will be very important. Should also include messages on a/i and effect of

broadcast messages.

From: thuel [thuel@dnrc.bell-labs.com]

Sent: Tuesday, January 25, 2000 3:53 PM

To: 'Tom LaPorta'

Cc: 'thuel@lucent.com'

Subject: RE: Comments on DHCP paper

Hi Tom,

I was wondering about you when I heard reports of LaGuardia being closed and long delays at

Newark and JFK. Sorry you had to postpone your trip but thanks a lot for reviewing the paper

and sending me your great comments. I've included my own comments below.

I'll make some changes as per your suggestions.

Thanks again for your input,
Sandy

>
> General
> -----
> 1. You use "should send" and like phrase when you should just state
> things like "the mobile sends" or the "the process is halted"

Good point.

>
> 2. Should we call the mobile host a mobile client instead (consistently
> throughout the paper)?

I don't have a strong preference either way so I had been sticking to the more conventional use

of the term mobile host, referring to the device rather than to a software client. Do you have a

reason to prefer the use of mobile client?

>
> 3. The main benefit of our approach is the we get all the benefits of
> DHCP (all options, re-use of servers, etc.). These must be spelled
> out very clearly or people will say "why assigning a 10* address, just
> assign the real thing).
>

I'll clarify this point.

> 4. We need a paragraph on how this would work with external FAs.
>

Luca and I were just talking about that this morning... Agreed.

> 5. Title: I think we should name our protocol. My proposal (off the
> top of my head): Transient Tunneling Protocol for Dynamic Addressing
> for Mobile Clients. The TTP could probably find other uses as well.
>

I was also thinking along the same lines but hadn't had time to really think about a name yet.

Transient Tunneling Protocol sounds like a great candidate. Sounds descriptive enough and has

a nice ring to it.

Thanks for the suggestion! I'll give it some more thought in case I can think

• of a better one

but this one will be hard to beat.

> Intro

> -----

> 1. Need strong statement of benefits/philosophy of our approach: minor
> changes to all protocols and servers, leverage existing stuff, get
> DHCP options
>

Yes, I had planned to beef up these arguments in the intro. on a second pass.
Will do.

> Section 4

> -----

> 1. Should figure 3 be broken into an "a" and "b" - one with a rely and
> one w/o?

I thought about this and decided it would take up too much space without
adding much value. Do

you think the text needs clarification regarding the non-relay case?

>

> Section 5

> -----

> 1. Could use a listing of our benefits

I deliberately kept these out of this section because I was planning to
emphasize these in the

intro. (as per your suggestion) first. We can then recall these benefits here,
if needed. Let

me take a crack at it and we'll reevaluate then.

> 2. Move 3rd paragraph into 5.1.

>

Note that the sentence refers to sections 5.1 and 5.2. I don't have a problem
in modifying it

accordingly but do you still think it should be moved?

> Section 5.1

> -----

> 1. We don't NEED to tear down the 10* tunnel - soft-state could expire.
> We should mention this as an option.

Good point, though I am not positive this would work. The reason is this:
we assume NAI-indexed listings of security associations at the HA. *IF* this

• implies that

tunneling entries must also be NAI-indexed, the temporary existence of multiple tunneling

entries for the same NAI can introduce resolution conflicts. Let me run this by Luca.

>

> 2. We should mention procedures with an FA here in 1 of 2 ways - either
> put a para here, or point to a para in the discussion section.

>

Let me think about what makes the most sense. My preference at this moment is to point to a

discussion elsewhere.

> 3. Does the B-bits cause any extra broadcast traffic to reach the
> mobile client/serving subnet? This may be very undesirable and may
> be a big drawback of our system.

>

Yes, this is an issue that we discussed a while back and then forgot about. Thanks for bringing

it up. Your concern is correct but it is possible to fix it by resetting the B bit with another

M-IP registration (after all DHCP transactions are over). I have to look at this again in more

detail to specify the message flow but the idea is to selectively enable broadcasting when DHCP

transactions are needed and disable it otherwise.

> Section 5.2

> -----

> 1. Is the first para just standard? If yes, say so.

>

Yes. Will do.

> Section 6.

> -----

> 1. Performance will be very important. Should also include messages
> on a/i and effect of broadcast messages.

Agreed. I'll be looking at this very closely.

From: Tom LaPorta [tlp@dnrc.bell-labs.com]

Sent: Tuesday, January 25, 2000 8:03 PM

To: thuel

Subject: Re: Comments on DHCP paper

> > 2. Should we call the mobile host a mobile client instead
> > (consistently throughout the paper)?
>
> I don't have a strong preference either way so I had been sticking to
> the more conventional use of the term mobile host, referring to the
> device rather than to a software client. Do you have a reason to
> prefer the use of mobile client?
>
Just a suggestion.

> > Section 4
> > -----
> > 1. Should figure 3 be broken into an "a" and "b" - one with a rely
> > and one w/o?
>
> I thought about this and decided it would take up too much space
> without adding much value. Do you think the text needs clarification
> regarding the non-relay case?
I think the text was fine. It's up to you.

>
> >
> > Section 5
> > -----
> > 1. Could use a listing of our benefits
>
> I deliberately kept these out of this section because I was planning
> to emphasize these in the intro. (as per your suggestion) first. We
> can then recall these benefits here, if needed. Let me take a crack
> at it and we'll reevaluate then.
OK.

>
> > 2. Move 3rd paragraph into 5.1.
> >
>
> Note that the sentence refers to sections 5.1 and 5.2. I don't have a
> problem in modifying it accordingly but do you still think it should be
> moved?
Not if it refers to both, but make that more clear.

>
> > Section 5.1
> > -----
> > 1. We don't NEED to tear down the 10* tunnel - soft-state could expire.
> > We should mention this as an option.
>

* > Good point, though I am not positive this would work. The reason is this:
> we assume NAI-indexed listings of security associations at the HA.
> *IF* this implies that tunneling entries must also be NAI-indexed, the
> temporary existence of multiple tunneling entries for the same NAI can
> introduce resolution conflicts. Let me run this by Luca.

OK
From: Luca Salgarelli [lsalgarelli@bell-labs.com]
Sent: Wednesday, January 26, 2000 10:10 AM
To: thuel
Subject: Re: draft-petri-mobileip-pipe-00.txt

Sandy,

> TLP> 1. We don't NEED to tear down the 10* tunnel - soft-state could expire.
> TLP> We should mention this as an option.
>
> SRT> Good point, though I am not positive this would work. The reason is
this:
> SRT> we assume NAI-indexed listings of security associations at the
> SRT> HA. *IF* this implies that tunneling entries must also be
> SRT> NAI-indexed, the temporary existence of multiple tunneling
> SRT> entries for the same NAI can introduce resolution conflicts. Let me run
this by Luca.
>
> Was this a valid concern or not?

I think it was. It's sort of complicated, because if we consider the NAI to be
the new

"indexing" element instead of the home-address, then it could be OK to follow
Tom's approach. In

this case, when the HA sees a new request, it would simply overwrite the
existing 10.* binding.

Nevertheless, this has not been written anywhere, and I suspect that it would
be very much

implementation dependent.

So, for the time being, I would stick with the sequential de-registration from
10.* followed by

the registration with the new addr.

Maybe we could just insert a note saying that alternative and more efficient
methods for

avoiding the explicit de-registration of the 10.* binding require more study.

Luca

•PS: thanks for the new version of the paper: I'll get you comments by tomorrow.

From: Luca Salgarelli [lsalgarelli@bell-labs.com]

Sent: Thursday, January 27, 2000 9:41 AM

To: Sandra Thuel

Subject: DHCP-MIP: comments on version 2

Hi Sandy.

So, here are my comments on version 2 of the draft.

1) I'm not sure the name TTP makes much sense: we are not introducing a protocol, it's just a

procedure. Switch it to Transient Tunneling Procedure for DHCP?

2) Probably you have already done this, but I think we should file for a patent before

)
publishing the paper.

2a) Another strong motivation for TTP, in addition to Tom's: while TTP is not needed for hosts

powering up in the home network (see point 4 below), power-ups on a foreign network is going to

be a much frequent case. E.g., use of MIP for corporate access.

3) Section 1, 2nd paragraph: why the colocated-COA should make any difference? I think making

this distinction here is asking for trouble.

After all, apart from the issues we were discussing yesterday, there are no high-level

differencies with TTP for the co-located and FA case.

4) Section 4, 2nd paragraph after figure 3: while it's true that unicasting may not work after a

move to a foreign network, it's also true that the client could use a reverse tunnel to get the

packet through, without defaulting to bcast. I don't think TTP is needed at all for clients

powering up at home and then moving into a foreign net.

5) Implementation: I used Linux for the client.

6) I was wondering if having separate sections for the FA case would make the paper more heavy.

Perhaps it is possible to present the cases side-by-side. After all, I don't see much difference

in TTP for the two cases.

Luca

1